

LISTING OF THE CLAIMS

What is claimed is:

1 1. (Previously Presented) In a communication device coupled to an
2 InfiniBand network and an external communication system, a method of
3 transferring a communication from the InfiniBand network to the external
4 communication system, the method comprising:

5 receiving packets for each of multiple queue pairs terminating at the
6 communication device;

7 for each of said queue pairs:

8 reassembling in a shared memory contents of said packets into
9 communications to be transmitted to the external communication system,
10 wherein the memory is a single contiguous memory structure shared by
11 multiple queue pairs and virtual lanes, and wherein a given packet payload
12 is stored contiguously in the memory; and

13 maintaining an associated linked list identifying locations in said
14 shared single contiguous memory in which said communications are
15 reassembled; and

16 when a communication is reassembled for a first queue pair, identifying to
17 a transmission module a portion of a first linked list associated with said first
18 queue pair, wherein said first linked list portion identifies shared single
19 contiguous memory locations in which said communication was reassembled.

1 2. (Original) The method of claim 1, further comprising:

2 only after said reassembled communication is transmitted, removing said
3 portion of said first linked list from said first linked list.

1 3. (Previously Presented) The method of claim 1, wherein said
2 reassembling comprises:

3 as said packets are received from the InfiniBand network, queuing said
4 contents directly into said shared single contiguous memory;

5 wherein said shared single contiguous memory serves as receive queues
6 for each of said multiple queue pairs.

1 4. (Previously Presented) The method of claim 3, further comprising:
2 transmitting said reassembled communication from said shared single
3 contiguous memory;

4 wherein said shared single contiguous memory serves as a transmit queue
5 for the external communication system.

1 5. (Original) The method of claim 1, wherein the external
2 communication system is an Ethernet network.

1 6. (Original) The method of claim 5, wherein said received packets
2 comprise portions of encapsulated Ethernet packets.

1 7. (Original) The method of claim 1, wherein said identifying to a
2 transmission module comprises transferring to the transmission module a set of
3 pointers identifying said reassembled communication rather than transferring to
4 the transmission module said reassembled communication.

1 8. (Original) The method of claim 1, wherein said identifying to a
2 transmission module comprises identifying to the transmission module:
3 a head of said portion of said first linked list; and

4 a tail of said portion of said first linked list.

1 9. (Previously Presented) The method of claim 1, wherein said
2 maintaining an associated linked list for said first queue pair comprises:

3 maintaining a head pointer configured to identify:

4 a first location in said shared single contiguous memory in which
5 contents of a first packet are stored; and

6 a first entry in a shared control structure, said first entry
7 corresponding to said first location in said shared single contiguous
8 memory; and

9 maintaining a tail pointer configured to identify:

10 a final location in said shared single contiguous memory in which
11 contents of a final packet are stored; and

12 a final entry in said shared control structure;

13 wherein each entry in said shared control structure that is part of said first
14 linked list, except for said final entry, identifies a subsequent entry in said shared
15 control structure and identifies a location in said shared single contiguous memory
16 corresponding to said subsequent entry.

1 10. (Original) The method of claim 9, further comprising:

2 protecting said shared control structure by one or more of:

3 separating bits of the structure to prevent double bit errors; and

4 providing single error correct and double error detect protection for one or
5 more entries in the control structure.

1 11. (Previously Presented) The method of claim 10, further
2 comprising: extending the single error correct and double error detect protection
3 to include said location in said shared single contiguous memory.

1 12. (Previously Presented) The method of claim 1, further comprising:
2 managing said linked lists for said queue pairs with a shared control;
3 wherein each said location in said shared single contiguous memory
4 corresponds to an entry in said shared control; and
5 wherein each entry in said shared control is configured to identify:
6 a subsequent entry within the same linked list; and
7 a location in said shared single contiguous memory corresponding
8 to said subsequent entry.

1 13. (Previously Presented) A computer readable medium storing
2 instructions that, when executed by a computer, cause the computer to perform a
3 method of transferring a communication from an InfiniBand network to a
4 communication system external to the InfiniBand network, the method
5 comprising:
6 for each of multiple queue pairs terminating at the communication device,
7 receiving packets;
8 for each of said queue pairs:
9 reassembling in a shared memory contents of said packets into
10 communications to be transmitted to the external communication system,
11 wherein the memory is a single contiguous memory structure shared by
12 multiple queue pairs and virtual lanes, and wherein a given packet payload
13 is stored contiguously in the memory; and
14 maintaining an associated linked list identifying locations in said
15 shared single contiguous memory in which said communications are
16 reassembled; and
17 when a communication is reassembled for a first queue pair, identifying to
18 a transmission module a portion of a first linked list associated with said first

19 queue pair, wherein said first linked list portion identifies shared single
20 contiguous memory locations in which said communication was reassembled

1 14. (Previously Presented) The computer readable medium of claim
2 13, the method further comprising:

3 as said packets are received from the InfiniBand network, queuing said
4 contents directly into said shared single contiguous memory, wherein said shared
5 single contiguous memory serves as receive queues for each of said multiple
6 queue pairs; and

7 transmitting said reassembled communication from said shared single
8 contiguous memory, wherein said shared single contiguous memory also serves as
9 a transmit queue for the external communication system.

1 15. (Original) The computer readable medium of claim 13, wherein
2 said identifying to a transmission module comprises transferring to the
3 transmission module a set of pointers identifying said reassembled
4 communication rather than transferring to the transmission module said
5 reassembled communication.

1 16. (Previously Presented) A method of storing a communication,
2 received from an InfiniBand network, for transmission external to the InfiniBand
3 network, the method comprising:

4 receiving a set of InfiniBand packets from an InfiniBand network, each
5 said packet comprising a portion of a communication to be transmitted external to
6 the InfiniBand network;

7 storing said communication portions in a single contiguous memory
8 shared among multiple queue pairs and virtual lanes of the InfiniBand network,
9 including a first queue pair through which said set of InfiniBand packets is

10 received, and wherein a given packet payload is stored contiguously in the
11 memory;

12 maintaining a first linked list for said first queue pair to identify locations
13 in said single contiguous memory in which said communication portions are
14 stored; and

15 when all of said communication portions are stored in said single
16 contiguous memory, scheduling said communication for transmission from said
17 single contiguous memory.

1 17. (Original) The method of claim 16, wherein said storing comprises
2 reassembling said communication portions into said communication.

1 18. (Previously Presented) The method of claim 16, wherein said
2 scheduling comprises:

3 identifying to a transmission module a first entry in said first linked list
4 corresponding to a location in said single contiguous memory in which a first
5 portion of said communication is stored; and

6 identifying to a transmission module a final entry in said first linked list
7 corresponding to a location in said single contiguous memory in which a final
8 portion of said communication is stored.

1 19. (Original) The method of claim 16, wherein said InfiniBand
2 packets comprise Send commands conveying said portions of said first
3 communication.

1 20. (Original) The method of claim 16, wherein said communication is
2 an Ethernet packet.

1 21. (Previously Presented) The method of claim 16, wherein said
2 maintaining a first linked list comprises:
3 maintaining a head pointer identifying:
4 a first entry in a control structure; and
5 a first location in said single contiguous memory in which a first
6 portion of said communication is stored; and
7 maintaining a tail pointer identifying:
8 a final entry in said control structure; and
9 a final location in said single contiguous memory in which a final
10 portion of said communication is stored;
11 wherein said first entry is linked to said final entry by zero or more
12 intermediate entries in said control structure, each said intermediate entry
13 corresponding to an intermediate location in said single contiguous memory in
14 which a portion of said communication is stored.

1 22. (Original) The method of claim 21, further comprising:
2 updating said first linked list to remove said first entry, said final entry and
3 said intermediate entries from said first linked list only after said communication
4 is transmitted.

1 23. (Previously Presented) The method of claim 16, wherein said
2 maintaining a first linked list comprises:
3 in a control structure, maintaining a first linked list of control entries,
4 wherein each of said control entries except a final control entry identifies:
5 a subsequent control entry; and
6 corresponding to said subsequent control entry, a location in said
7 single contiguous memory in which data received through said first queue
8 pair are stored.

1 24. (Previously Presented) The method of claim 23, wherein said
2 control structure and said single contiguous memory are separate physical memory
3 structures.

1 25. (Previously Presented) The method of claim 24, wherein said
2 control structure and said single contiguous memory employ common addressing
3 via said first linked list.

1 26. (Previously Presented) The method of claim 23, wherein said
2 maintaining further comprises:
3 maintaining a head pointer identifying a first control entry in said first
4 linked list and a first location in said single contiguous memory; and
5 maintaining a tail pointer identifying said final control entry in said first
6 linked list and a final location in said single contiguous memory.

1 27. (Previously Presented) The method of claim 23, further
2 comprising:
3 identifying to a transmission module a sub-list of said first linked list of
4 control entries, wherein said sub-list includes control entries corresponding to all
5 locations in said single contiguous memory in which portions of said
6 communication are stored.

1 28. (Original) The method of claim 27, further comprising:
2 removing said sub-list of control entries from said first linked list only
3 after said communication is transmitted.

1 29. (Previously Presented) A computer readable medium storing

2 instructions that, when executed by a computer, cause the computer to perform a
3 method of storing a communication, received from an InfiniBand network, for
4 transmission external to the InfiniBand network, the method comprising:

5 receiving a set of InfiniBand packets from an InfiniBand network, each
6 said packet comprising a portion of a communication to be transmitted external to
7 the InfiniBand network;

8 storing said communication portions in a single contiguous memory
9 shared among multiple queue pairs and virtual lanes of the InfiniBand network,
10 including a first queue pair through which said set of InfiniBand packets is
11 received, wherein a given packet payloads is stored contiguously in the memory;

12 maintaining a first linked list for said first queue pair to identify locations
13 in said single contiguous memory in which said communication portions are
14 stored; and

15 when all of said communication portions are stored in said single
16 contiguous memory, scheduling said communication for transmission from said
17 single contiguous memory.

1 30. (Previously Presented) The computer readable medium of claim
2 29, wherein said scheduling comprises:

3 identifying to a transmission module a first entry in said first linked list
4 corresponding to a location in said single contiguous memory in which a first
5 portion of said communication is stored; and

6 identifying to a transmission module a final entry in said first linked list
7 corresponding to a location in said single contiguous memory in which a final
8 portion of said communication is stored.

1 31. (Previously Presented) The computer readable medium of claim
2 29, wherein said maintaining a first linked list comprises:

3 in a control structure, maintaining a first linked list of control entries,
4 wherein each of said control entries except a final control entry identifies:
5 a subsequent control entry; and
6 corresponding to said subsequent control entry, a location in said
7 single contiguous memory in which data received through said first queue
8 pair are stored;
9 maintaining a head pointer identifying a first control entry in said first
10 linked list and a first location in said single contiguous memory; and
11 maintaining a tail pointer identifying said final control entry in said first
12 linked list and a final location in said single contiguous memory.

1 32. (Previously Presented) The computer readable medium of claim
2 31, wherein the method further comprises:
3 identifying to a transmission module a sub-list of said first linked list of
4 control entries, wherein said sub-list includes control entries corresponding to all
5 locations in said single contiguous memory in which portions of said
6 communication are stored.

1 33. (Original) The computer readable medium of claim 32, wherein the
2 method further comprises:
3 removing said sub-list of control entries from said first linked list only
4 after said communication is transmitted.

1 34. (Previously Presented) A computer readable memory structure
2 configured for simultaneously queuing contents of packets as they are received
3 from an InfiniBand network and storing communications reassembled from said
4 contents for transmission external to the InfiniBand network, the memory
5 structure comprising:

6 a single contiguous memory shared by multiple queue pairs and virtual
7 lanes, comprising:

8 for each of a plurality of InfiniBand communication connections,
9 memory buckets for storing contents of packets received on the
10 communication connection, wherein a given packet payload is stored
11 contiguously in the memory within and across the buckets;

12 wherein said contents of said packets are reassembled into
13 communications in said memory buckets as said contents are stored;
14 a shared control, comprising:

15 for each of said communication connections, a linked list of control
16 entries, wherein each said control entry is configured to identify:

17 a subsequent control entry in said linked list; and
18 in said shared single contiguous memory, a memory bucket
19 corresponding to said subsequent control entry;

20 wherein said shared single contiguous memory and said shared control are
21 configured for access by:

22 an InfiniBand receive module configured to receive said packets;
23 and

24 a transmit module configured to transmit said communications
25 external to the InfiniBand network.

1 35. (Previously Presented) The computer readable memory structure
2 of claim 34, wherein said shared control further comprises:

3 a linked list of free entries, wherein each said free entry is configured to
4 identify, in said shared single contiguous memory, a free memory bucket.

1 36. (Previously Presented) The computer readable memory structure of
2 claim 34, wherein said shared control is protected by one or more of:

3 separating bits of the structure to prevent double bit errors; and
4 providing single error correct and double error detect protection for one or
5 more control entries in the shared control.

1 37. (Previously Presented) The computer readable memory structure of
2 claim 36, wherein a control entry in said shared control is further protected by
3 extending the single error correct and double error detect protection to include the
4 identity of a memory bucket in said shared single contiguous memory.

1 38. (Previously Presented) An apparatus for storing data received from
2 an InfiniBand network, for transfer to an external communication system,
3 comprising:

4 an InfiniBand receive module configured to receive packets from a
5 plurality of InfiniBand communication connections;
6 a transmit module configured to transmit communications to a
7 communication system external to the InfiniBand network via one or more
8 outbound ports;

9 a single contiguous memory shared between said InfiniBand receive
10 module and said transmit module, wherein a given packet payload is stored
11 contiguously in the memory; and

12 a control, shared between said InfiniBand receive module and said
13 transmit module, for maintaining a linked list for each of said communication
14 connections and for each of the outbound ports.

1 39. (Original) The apparatus of claim 38, further comprising, for each
2 said communication connection:

3 a head pointer identifying a head of said linked list for said communication
4 connection; and

5 a tail pointer identifying a tail of said linked list for said communication
6 connection.

1 40. (Original) The apparatus of claim 38, further comprising, for each
2 of the outbound ports:

3 a head pointer identifying a head of said linked list for the outbound port;
4 and

5 a tail pointer identifying a tail of said linked list for the outbound port.

1 41. (Original) The apparatus of claim 40, further comprising, for each
2 outbound queue of each of the outbound ports:

3 a head pointer identifying a head of said linked list for the outbound
4 queue; and

5 a tail pointer identifying a tail of said linked list for the outbound queue.

1 42. (Original) The apparatus of claim 38, wherein said communication
2 connections are queue pairs.

1 43. (Original) The apparatus of claim 38, wherein said communication
2 connections are virtual lanes.

1 44. (Original) The apparatus of claim 38, wherein the external
2 communication system comprises an Ethernet network.

1 45. (Original) The apparatus of claim 38, wherein the external
2 communication system comprises a SONET network.